29.1 Power Distribution Measurements of the Dual Core PowerPC™ 970MP Microprocessor

Hendrik F. Hamann[1], Alan Weger[1], James Lacey[1], Erwin Cohen[2], Craig Atherton[2]

[1]IBM T.J. Watson Research Center, Yorktown Heights, NY
[2]IBM Systems & Technology Group, Essex Junction, VT

Power dissipation has emerged as a major design constraint limiting performance, reliability and manufacturing yields of today's microprocessors. The precise impact of this power limit is determined by the details of the power *distributions* within the microprocessor circuit. However, these power distributions are typically inferred from model calculations with numerous underlying assumptions. In order to understand power distributions and the impact of hot spots we have developed an experimental method, which not only allows for real-time imaging of temperature distributions at a resolution far beyond the capability of on-chip thermal sensors but it also can determine detailed power maps of microprocessors under a wide variety of operating conditions. In this paper we apply Spatially-resolved Imaging of Microprocessor Power (SIMP) to measure the detailed power distribution of the dual core PowerPC™ 970MP microprocessor as a function of frequency, voltage and workload. The data shows a significant movement of the hot spot for one and two core operations, respectively. Finally, we compare the extent of the hot spots between two competitive microprocessors under the same workload.

Figure 29.1.1 illustrates the experimental setup of the SIMP method, which involves two steps [1-2]. In the first step we use infra-red (IR) thermal imaging to measure the thermal distribution of the microprocessor. Specifically, an IR-transparent liquid flows directly through a micro-duct of a specifically designed cooling cell over the microprocessor while the IR radiation from the backside of the processor is projected onto an InSb array detector. The spatial resolution of the IR camera is ~ 60 microns. The heat removal rate of this transparent heat sink can be tuned up to 200 W/cm^2 by the flow speed of the liquid to mimic various cooling conditions. The backside of the microprocessor is coated with a thin blackbody coating. In order to limit the spreading of the heat, the microprocessor die is thinned to ~ 100 microns thickness. Figure 29.1.1 also shows a photograph of the actual experiment setup, where the IR camera and the cooling cell are mounted onto a Vanguard microprocessor tester.

The second step of the SIMP technique is illustrated in Fig. 29.1.2, where we represent the measured chip temperature as a superposition of individual temperature fields for each power source on the chip. Consequently, we can relate the chip power to the chip temperature at each unit cell by $\underline{A}\ \underline{P} = \underline{T}$, where $\underline{A}$ is the thermal resistance matrix, $\underline{P}$ the power at the circuitry level and $\underline{T}$ the corresponding (measured) temperature vector at the die backside. The coefficients of the A-matrix are thermal resistances (e.g., units of K/W). For example and referring to Fig. 29.1.2, $a_{31,52}$ is the temperature increase in cell 31 if one unit of power is applied in cell 52. The thermal resistance matrix can be either measured (by applying power sources systematically across the chip (e.g., by using a focused laser beam) and measuring the corresponding temperature fields) or calculated by using fluid dynamic model simulations. In this study we employed calculated matrixes (30×30), which were validated using thermal test chips. The spatial resolution of the power maps shown in this study is 0.4×0.4mm^2, which is sufficient for packaging and accurately calculating the temperature distribution.

Figure 29.1.3 shows various measured temperature and subsequently derived power distributions for the PowerPC™ 970MP microprocessor obtained with SIMP method. Each temperature / power column pair in Fig. 29.1.3 refers to a different color bar. In the first two columns, the respective temperature and power distributions are shown for three different workloads: DC, IDLE, and a high power AVP (here for 1.6GHz, Vdd=1.2V). In the DC power map the receiver terminators (located at the right lower corner of the microprocessor) are clearly visible. Although the actual temperature impact of these devices is quite low their footprint can be clearly observed in each of the power maps of the subsequent images in Fig. 29.1.3. The data demonstrates how well the SIMP technique de-convolves the measured temperature images. In the third and fourth column, the temperature and corresponding power maps for the IDLE workload are shown (here for 2.0GHz and Vdd=1.2V). In the first row both cores are running and in the second row the left core is running while the right core has been entirely turned off. Finally, in the third row the left core has been powered down. In the fifth and six column, the respective temperature and power images are displayed for a high power AVP (here for 1.6GHz and Vdd=1.15V). In the first row it runs on both cores while in the second row it only runs on the left core. In summary, the various images of Fig. 29.1.3 demonstrate how well the SIMP technique picks up the changes in power distributions for the different conditions and workloads.

In Fig. 29.1.4 we show the temperature distributions for a high power AVP running on one and two cores, respectively. As it is very evident from Fig. 29.1.4 the hot spot moves significantly (by ~ 2mm) between the one and two core cases. Most microprocessors have thermal sensors on the chip, which are used to manage chip temperatures. However, this example shows clearly that the differences between the chip peak temperatures and the thermal sensor readings are workload dependent, which can have naturally wide implications for power and thermal management schemes.

Figure 29.1.5 allows a closer look at how the power is actually distributed over the different units (i.e., here for the high power AVP at 1.6GHz and Vdd=1.2V). The alignment of the experimental power data with the floor plan is quite remarkable considering the coarse granularity of the power map. The highest power density regions are sub units of the vector engine, the FXUs and sub units of the ISU. This power map is useful to explain why the hot spot shifts as shown in Fig. 29.1.4. With both cores running, the hotspot is governed by the thermal cross-talk of the FXUs. However, if only one core runs the vector engine sub-unit dominates and thus the hotspot is shifting towards the upper corner. By comparing the measured power maps with modeling predictions it is now possible to improve (pre-silicon) power model tools. This will be critical in order to consider more accurately the resulting power maps of microprocessor circuits at an early stage of the design cycle.

Finally, we investigate how the impact of the hot spot of the PowerPC™ 970MP compares to a competitive microprocessor. Specifically, we compare in Fig. 29.1.6 the temperature distributions of the predecessor of the PowerPC™ 970MP chip (single core, 90nm, 2.5GHz) *with* a competitive microprocessor (3.2GHz, 90nm). Both chips are observed under identical cooling conditions (e.g., the same flow rate and duct height), they run the same workload (GCC) at almost identical performances and neither chip is thinned. The PowerPC™ 970MP predecessor is using 79W while the competitive chip runs at 85W. The images are shown to scale. Although the predecessor chip has the higher *average* power density (because it is a smaller chip), it certainly shows less of a hot spot (i.e., lower temperatures for identical workload, performance and cooling solution) than its competitor.

References:

[1] H.F. Hamann et al., "Spatially Resolved Imaging of Microprocessor Power," *ISSCC Dig. Tech. Papers*, pp. 16, Feb., 2005.
[2] S. Guha et al.,"Transparent Cooling Duct", *US Patent application* 20050094706, 2005.

 1-4244-0079-1/06/$20.00 ©2006 IEEE

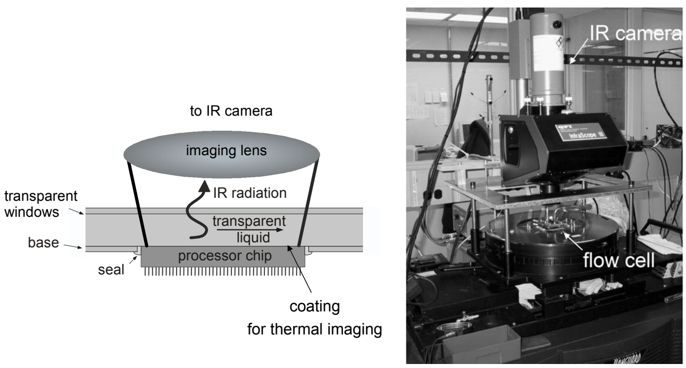

Figure 29.1.1: Schematic of the cooling cell and photograph of the experimental setup.

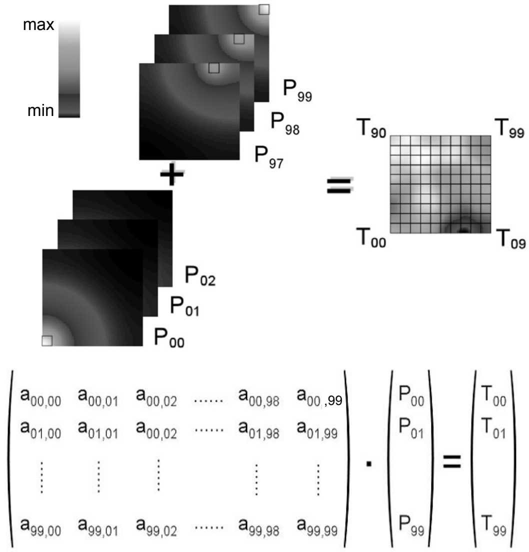

Figure 29.1.2: Illustration of the conversion from temperature to power distribution.

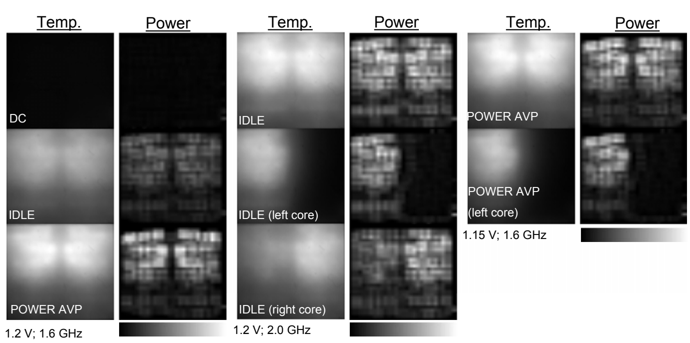

Figure 29.1.3: Power and temperature distribution of the PowerPC™ 970MP microprocessor under various conditions and workloads. The die size (including kerfs) is 13.193×11.595 mm^2.

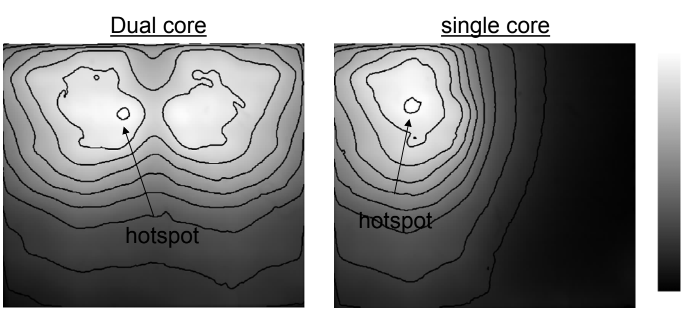

Figure 29.1.4: Hotspot movement for dual and single core operation.

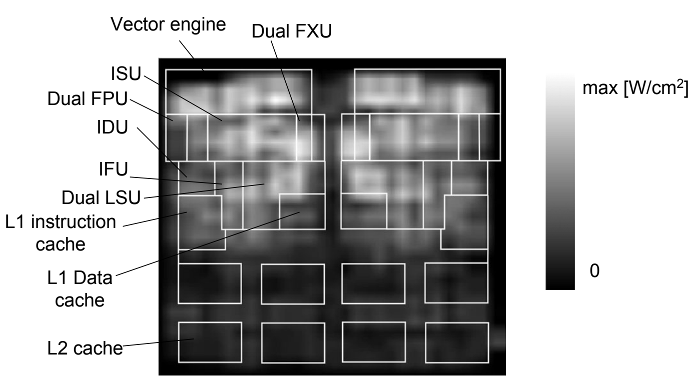

Figure 29.1.5: Overlay of PowerPC™ 970MPpower map with the floor plan for the high power AVP workload.

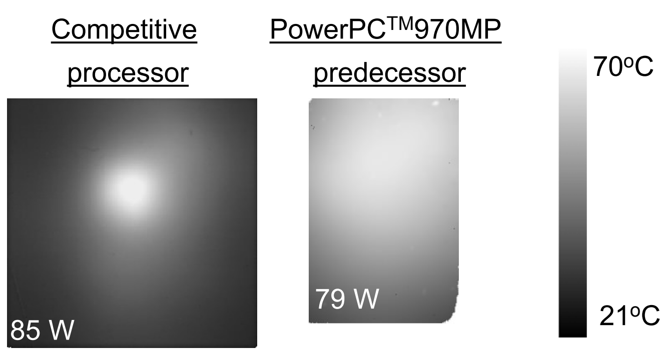

Figure 29.1.6: Comparison of the thermal distribution of the PowerPC™ 970MP predecessor with a competitive microprocessor running the same workload under the same cooling conditions.

29